arXiv:1402.4037v2 [cs.DS] 18 Feb 2015

Near-Linear Query Complexity for Graph Inference

Sampath Kannan[1], Claire Mathieu[2], and Hang Zhou[2]

[1] Department of Computer and Information Science
University of Pennsylvania, Philadelphia, PA, USA.
kannan@cis.upenn.edu
[2] Département d'Informatique UMR CNRS 8548,
École Normale Supérieure, Paris, France
{cmathieu,hangzhou}@di.ens.fr

Abstract. How efficiently can we find an unknown graph using distance or shortest path queries between its vertices? Let $G = (V, E)$ be an unweighted, connected graph of bounded degree. The edge set E is initially unknown, and the graph can be accessed using a *distance oracle*, which receives a pair of vertices (u, v) and returns the distance between u and v. In the *verification* problem, we are given a hypothetical graph $\hat{G} = (V, \hat{E})$ and want to check whether G is equal to $\hat{G}$. We analyze a natural greedy algorithm and prove that it uses $n^{1+o(1)}$ distance queries. In the more difficult *reconstruction* problem, $\hat{G}$ is not given, and the goal is to find the graph G. If the graph can be accessed using a *shortest path oracle*, which returns not just the distance but an actual shortest path between u and v, we show that extending the idea of greedy gives a reconstruction algorithm that uses $n^{1+o(1)}$ shortest path queries. When the graph has bounded treewidth, we further bound the query complexity of the greedy algorithms for both problems by $\tilde{O}(n)$. When the graph is chordal, we provide a randomized algorithm for reconstruction using $\tilde{O}(n)$ distance queries.

1 Introduction

How efficiently can we find an unknown graph using distance or shortest path queries between its vertices? This is a natural theoretical question from the standpoint of recovery of hidden information. This question is related to the *reconstruction* of Internet networks. Discovering the topology of the Internet is a crucial step for building accurate network models and designing efficient algorithms for Internet applications. Yet, this topology can be extremely difficult to find, due to the dynamic structure of the network and to the lack of centralized control. The network reconstruction problem has been studied extensively [1,2,6,7,11,16]. Sometimes we have some idea of what the network should be like, based perhaps on its state at some past time, and we want to check whether our image of the network is correct. This is network *verification* and has received attention recently [2,4,7]. This is an important task for routing,

error detection, or ensuring service-level agreement (SLA) compliance, etc. For example, Internet service providers (ISPs) offer their customers services that require quality of service (QoS) guarantees, such as voice over IP services, and thus need to check regularly whether the networks are correct.

The topology of Internet networks can be investigated at the router and autonomous system (AS) level, where the set of routers (ASs) and their physical connections (peering relations) are the vertices and edges of a graph, respectively. Traditionally, we use tools such as traceroute and mtrace to infer the network topology. These tools generate path information between a pair of vertices. It is a common and reasonably accurate assumption that the generated path is the shortest one, i.e., minimizes the hop distance between that pair. In our first theoretical model, we assume that we have access to any pair of vertices and get in return their shortest path in the graph. Sometimes routers block traceroute and mtrace requests (e.g., due to privacy and security concerns), thus the inference of topology can only rely on delay information. In our second theoretical model, we assume that we get in return the hop distance between a pair of vertices. The second model was introduced by Mathieu and Zhou [11].

Graph inference using queries that reveal partial information has been studied extensively in different contexts, independently stemming from a number of applications. Beerliova et al. [2] studied network verification and reconstruction using an oracle, which, upon receiving a node q, returns all shortest paths from q to all other nodes, instead of one shortest path between a pair of nodes as in our first model. Erlebach et al. [7] studied network verification and reconstruction using an oracle which, upon receiving a node q, returns the distances from q to all other nodes in the graph, instead of the distance between a pair of nodes as in our second model. They showed that minimizing the number of queries for verification is NP-hard and admits an $O(\log n)$-approximation algorithm. In the *network realization* problem, we are given distances between certain pairs of vertices and asked to determine the sparsest graph (in the unweighted case) or the graph of least total weight that realizes these distances. This problem was shown to be NP-hard [5]. In evolutionary biology, the goal is to reconstruct evolutionary trees, thus the hidden graph has a tree structure. See for example [8,10,15]. One may query a pair of species and get in return the distance between them in the (unknown) tree. In our reconstruction problem, we allow the hidden graph to have an arbitrary connected topology, not necessarily a tree structure.

1.1 The Problem

Let $G = (V, E)$ be a hidden graph that is connected and unweighted, where $|V| = n$. We consider two query oracles. A *shortest path oracle* receives a pair $(u, v) \in V^2$ and returns a shortest path between u and v. A *distance oracle* receives a pair $(u, v) \in V^2$ and returns the number of edges on a shortest path between u and v.

In the *graph reconstruction* problem, we are given the vertex set V and have access to either a distance oracle or a shortest path oracle. The goal is to find every edge in E.

Table 1: Results (for bounded degree graphs). New results are in bold.

Objective	*Query complexity*
verification using a distance oracle	$\boldsymbol{n^{1+o(1)}}$ **bounded treewidth:** $\tilde{O}(n)$ **(Thm 1 and Thm 3)**
reconstruction using a shortest path oracle	
reconstruction using a distance oracle	$\tilde{O}(n^{3/2})$ [11] $\boldsymbol{\Omega(n \log n / \log\log n)}$ **(Thm 5)** outerplanar: $\tilde{O}(n)$ [11] **chordal:** $\boldsymbol{\tilde{O}(n)}$ **(Thm 4)**

In the *graph verification* problem, again we are given V and have access to one of the two oracles. In addition, we are given an unweighted, connected graph $\hat{G} = (V, \hat{E})$. The goal is to check whether $\hat{G}$ is correct, that is, whether $\hat{G} = G$.

The efficiency of an algorithm is measured by its *query complexity*[3], i.e., the number of queries to an oracle. We focus on query complexity, while all our algorithms are of polynomial time and space. We note that $O(n^2)$ queries are enough for both reconstruction and verification using a distance oracle or a shortest path oracle: we only need to query every pair of vertices.

Let Δ denote the maximum degree of any vertex in the graph G. Unless otherwise stated, we assume that Δ is bounded, which is reasonable for real networks that we want to reconstruct or verify. Indeed, when Δ is $\Omega(n)$, both reconstruction and verification require $\Omega(n^2)$ distance or shortest path queries, see Section 6.1.[4]

Let us focus on bounded degree graphs. It is not hard to see that $\Omega(n)$ queries are required. The central question in this line of work is therefore: **Is the query complexity linear, quadratic, or somewhere in between?** In [11], Mathieu and Zhou provide a first answer: the query complexity for reconstruction using a distance oracle is subquadratic: $\tilde{O}(n^{3/2})$. In this paper, we show that the query complexity for reconstruction using a shortest path oracle or verification using a distance oracle is near-linear: $n^{1+o(1)}$.

1.2 Our Results

Verification.

Theorem 1. *For graph verification using a distance oracle, there is a deterministic algorithm (Algorithm 1) with query complexity* $n^{1+O\left(\sqrt{(\log\log n + \log\Delta)/\log n}\right)}$, *which is* $n^{1+o(1)}$ *when the maximum degree* $\Delta = n^{o(1)}$. *If the graph has treewidth* tw, *the query complexity can be further bounded by* $O(\Delta(\Delta + \mathrm{tw}\log n)n\log^2 n)$, *which is* $\tilde{O}(n)$ *when* Δ *and* tw *are* $O(\mathrm{polylog}\, n)$.

[3] Expected query complexity in the case of randomized algorithms.

[4] We note that the $\Omega(n^2)$ lower bound holds even when the graph is restricted to chordal or to bounded treewidth.

The main task for verification is to confirm the *non-edges* of the graph. Algorithm 1 is greedy: every time it makes a query that confirms the largest number of non-edges that are not yet confirmed. To analyze the algorithm, first, we show that its query complexity is $O(\log n)$ times the optimal number of queries OPT for verification. This is based on a reduction to the Set-Cover problem, see Section 3.1. It only remains to bound OPT.

To bound OPT and get the first statement in Theorem 1, it is enough to prove the desired bound for a different verification algorithm. This algorithm is a more sophisticated recursive version of the algorithm in [11]. Recursion is a challenge because, when we query the pair (u, v) in a recursive subgraph, the oracle returns the distance between u and v in the entire graph, not just within the subgraph. Thus new ideas are introduced for the algorithmic design. See Section 3.3.

To show the second statement in Theorem 1, similarly, we design another recursive verification algorithm with query complexity $\tilde{O}(n)$ for graphs of bounded treewidth. The algorithm uses some bag of a tree decomposition to separate the graph into balanced subgraphs, and then recursively verifies each subgraph. The same obstacle to recursion occurs. Our approach here is to add a few weighted edges to each subgraph in order to preserve the distance metric. See Section 3.4.

We note that each query to a distance oracle can be simulated by the same query to a shortest path oracle. So from Theorem 1, we have:

Corollary 2. *For graph verification using a shortest path oracle, Algorithm 1 achieves the same query complexity as in Theorem 1.*

Reconstruction.

Theorem 3. *For graph reconstruction using a shortest path oracle, there is a deterministic algorithm (Algorithm 5) that achieves the same query complexity as in Theorem 1.*

The key is to formulate this problem as a problem of verification using a distance oracle, so that we get the same query complexity as in Theorem 1. We extend ideas of the greedy algorithm in Theorem 1 to design Algorithm 5, and we show that each query to a shortest path oracle makes as much progress for reconstruction as the corresponding query to a distance oracle would have made for verifying a given graph. The main realization here is that reconstruction can be viewed as the verification of a dynamically changing graph. See Section 4.

Theorem 4. *For reconstruction of* **chordal graphs** *using a distance oracle, there is a randomized algorithm (Algorithm 8) with query complexity* $O\big(\Delta^3 2^\Delta \cdot n(2^\Delta + \log^2 n)\log n\big)$*, which is* $\tilde{O}(n)$ *when the maximum degree* Δ *is* $O(\log\log n)$.

The algorithm first finds a separator using random sampling and statistical estimates, as in [11]. Then it partitions the graph into subgraphs with respect to this separator and recurses on each subgraph. However, the separator here is a clique instead of an edge in [11] for outerplanar graphs. Thus the main

difficulty is to design and analyze a more general tool for partitioning the graph, see Section 5.1. The reconstruction algorithm is in Section 5.3.

On the other hand, graph reconstruction using a distance oracle has a lower bound that is slightly higher than trivial $\Omega(n)$ bound, as in the following theorem. Its proof is in Section 6.2.

Theorem 5. *For graph reconstruction using a distance oracle, assuming the maximum degree $\Delta \geq 3$ is such that $\Delta = o\left(n^{1/2}\right)$, any algorithm has query complexity $\Omega(\Delta n \log n / \log\log n)$.*

It is an outstanding open question whether there is a reconstruction algorithm using a near-linear number of queries to a distance oracle for degree bounded graphs in general.

2 Notation and Preliminaries

Let δ be the distance metric of G. For a subset of vertices $S \subseteq V$ and a vertex $v \in V$, define $\delta(S, v)$ to be $\min_{s \in S} \delta(s, v)$. For $v \in V$, let $N(v) = \{u \in V : \delta(u, v) \leq 1\}$ and let $N_2(v) = \{u \in V : \delta(u, v) \leq 2\}$. For $S \subseteq V$, let $N(S) = \bigcup_{s \in S} N(s)$. We define $\hat{\delta}$, $\hat{N}$, and $\hat{N}_2$ similarly with respect to the graph $\hat{G}$.

A pair of vertices $\{u, v\} \subseteq V$ is called a *non-edge* of the graph $G = (V, E)$ if $\{u, v\} \notin E$.

For a subset of vertices $S \subseteq V$, let $G[S]$ be the subgraph induced by S. For a subset of edges $H \subseteq E$, we identify H with the subgraph induced by the edges of H. Let δ_H denote the distance metric of the subgraph H.

For a vertex $s \in V$ and a subset $T \subseteq V$, define $\textsc{Query}(s, T)$ as $\textsc{Query}(s, t)$ for every $t \in T$. For subsets $S, T \subseteq V$, define $\textsc{Query}(S, T)$ as $\textsc{Query}(s, t)$ for every $(s, t) \in S \times T$.

Definition 6. *A subset $S \subseteq V$ is a* β-balanced separator *of the graph $G = (V, E)$ (for $\beta < 1$) if the size of every connected component of $G \setminus S$ is at most $\beta|V|$.*

Definition 7. *A* tree decomposition *of a graph $G = (V, E)$ is a tree T with nodes $n_1, n_2, \ldots, n_\ell$. Node n_i is identified with a* bag *$S_i \subseteq V$, satisfying the following conditions:*

1. *For every vertex v in G, the nodes whose bags contain v form a connected subtree of T.*
2. *For every edge (u, v) in G, some bag contains both u and v.*

The width *of the decomposition is the size of the largest bag minus 1, and the* treewidth *of G is the minimum width over all possible tree decompositions of G.*

Lemma 8 ([13]). *Let G be a graph of treewidth k. Any tree decomposition of width k contains a bag S that is a $(1/2)$-balanced separator of G.*

A graph is *chordal* if every cycle of length greater than three has a chord: namely, an edge connecting two nonconsecutive vertices on the cycle. An introduction to chordal graphs can be found in *e.g.,* [3].

Lemma 9 ([3]). *Let G be a chordal graph. Then G has a tree decomposition where every bag is a maximal clique*[5] *and every maximal clique appears exactly once in this decomposition.*

From Lemmas 8 and 9, we have:

Corollary 10. *Let G be a chordal graph of maximum degree Δ. Then G has treewidth at most Δ, and there exists a clique $S \subseteq V$ of size at most $\Delta + 1$ that is a $(1/2)$-balanced separator of G.*

3 Proof of Theorem 1

3.1 Greedy Algorithm

The task of verification comprises verifying that every edge in $\hat{G}$ is an edge in G, and verifying that every non-edge of $\hat{G}$ is a non-edge of G. The second part, called *non-edge verification*, is the main task for graphs of bounded degree.[6]

Theorem 11. *For graph verification using a distance oracle, there is a deterministic greedy algorithm (Algorithm 1) that uses at most $\Delta n + (\ln n + 1) \cdot OPT$ queries, where OPT is the optimal number of queries for non-edge verification.*

Now we prove Theorem 11. Let $\widehat{NE}$ be the set of the non-edges of $\hat{G}$. For each pair of vertices $(u, v) \in V^2$, we define $S_{u,v} \subseteq \widehat{NE}$ as follows:

$$S_{u,v} = \left\{ \{a, b\} \in \widehat{NE} : \hat{\delta}(u, a) + \hat{\delta}(b, v) + 1 < \hat{\delta}(u, v) \right\}. \tag{1}$$

The following two lemmas relate the sets $S_{u,v}$ with non-edge verification.

Lemma 12. *Assume that $\hat{E} \subseteq E$. Let $(u, v) \in V^2$ be such that $\delta(u, v) = \hat{\delta}(u, v)$. Then every pair $\{a, b\} \in S_{u,v}$ is a non-edge of G.*

Proof. Let $\{a, b\}$ be any pair in $S_{u,v}$. By the triangle inequality, $\delta(u, a) + \delta(a, b) + \delta(b, v) \geq \delta(u, v) = \hat{\delta}(u, v)$. By the definition of $S_{u,v}$ and using $\hat{E} \subseteq E$, we have $\hat{\delta}(u, v) > \hat{\delta}(u, a) + \hat{\delta}(b, v) + 1 \geq \delta(u, a) + \delta(b, v) + 1$. Thus $\delta(a, b) > 1$, i.e., $\{a, b\}$ is a non-edge of G. □

Lemma 13. *If a set of queries T verifies that every non-edge of $\hat{G}$ is a non-edge of G, then $\bigcup_{(u,v)\in T} S_{u,v} = \widehat{NE}$.*

[5] A maximal clique is a clique which is not contained in any other clique.

[6] In non-edge verification, we always assume that $\hat{E} \subseteq E$.

Proof. Assume, for a contradiction, that some $\{a, b\} \in \widehat{NE}$ does not belong to any $S_{u,v}$ for $(u, v) \in T$. Consider adding $\{a, b\}$ to the set of edges of $\hat{E}$: this will not create a shorter path between u and v, for any $(u, v) \in T$. Thus including $\{a, b\}$ in $\hat{E}$ is consistent with the answers of all queries in T. This contradicts the assumption that T verifies that $\{a, b\}$ is a non-edge of G. □

From Lemmas 12 and 13, the non-edge verification is equivalent to the Set-Cover problem with the universe $\widehat{NE}$ and the sets $\{S_{u,v} : (u, v) \in V^2\}$. The Set-Cover instance can be solved using the well-known greedy algorithm [9], which gives a $(\ln n + 1)$-approximation. Hence our greedy algorithm for verification (Algorithm 1). For the query complexity, first, verifying that $\hat{E} \subseteq E$ takes at most Δn queries, since the graph has maximum degree Δ. The part of non-edge verification uses a number of queries that is at most $(\ln n + 1)$ times the optimal number of queries. This proves Theorem 11.

Algorithm 1 Greedy Verification

```
1: procedure VERIFY(Ĝ)
2:     for {u, v} ∈ Ê do QUERY(u, v)
3:     if some {u, v} ∈ Ê has δ(u, v) ≠ δ̂(u, v) then return no
4:     Y ← ∅
5:     while Ê ∪ Y does not cover all vertex pairs do
6:         choose (u, v) that maximizes |S_{u,v} \ Y|        ▷ S_{u,v} defined in Equation (1)
7:         Query(u, v)
8:         if δ(u, v) = δ̂(u, v) then
9:             Y ← Y ∪ S_{u,v}
10:        else
11:            return no
12:    return yes
```

3.2 Bounding *OPT* to Prove Theorem 1

From Theorems 11, in order to prove Theorem 1, we only need to bound OPT, as in the following two theorems.

Theorem 14. *For graph verification using a distance oracle, the optimal number of queries OPT for non-edge verification is $n^{1+O\left(\sqrt{(\log\log n+\log\Delta)/\log n}\right)}$.*

Theorem 15. *For graph verification using a distance oracle, if the graph has treewidth* tw*, then the optimal number of queries OPT for non-edge verification is $O(\Delta(\Delta + \mathrm{tw}\log n)n\log n)$.*

Theorem 1 follows trivially from Theorems 11, 14, and 15, by noting that both Δ and $\log n$ are smaller than $n^{\sqrt{(\log\log n+\log\Delta)/\log n}}$. The proof of Theorem 14 is in Section 3.3, and the proof of Theorem 15 is in Section 3.4.

3.3 Proof of Theorem 14

To show Theorem 14, we provide a recursive algorithm for non-edge verification with the query complexity in the theorem statement. As in [11], the algorithm selects a set of *centers* partitioning V into Voronoi cells and expands them slightly so as to cover all edges of G. But unlike [11], instead of using exhaustive search inside each cell, the algorithm verifies each cell recursively. The recursion is a challenge because the distance oracle returns the distance in the entire graph, not in the cell. Straightforward attempts to use recursion lead either to subcells that do not cover every edge of the cell, or to excessively large subcells. To make the recursion work, we allow selection of centers *outside* the cell, while still limiting the subcells to being contained *inside* the cell (Figure 1). This simple but subtle setup is one novelty of the algorithmic design.

Let $U \subseteq V$ represents the set of vertices for which we are currently verifying the induced subgraph. The goal is to verify that every non-edge of $\hat{G}[U]$ is a non-edge of $G[U]$. This is equivalent to verifying that every edge of $G[U]$ is an edge of $\hat{G}[U]$.

The algorithm uses a subroutine to find *centers* $A \subseteq V$ such that the vertices of U are roughly equipartitioned into the Voronoi cells centered at vertices in A. For a set of centers $A \subseteq V$ and a vertex $w \in V$, let $\hat{C}_A(w) = \{v \in V : \hat{\delta}(w, v) < \hat{\delta}(A, v)\}$, which represents the Voronoi cell of w if w is added to the set of centers. We note that $\hat{C}_A(w) = \emptyset$ for $w \in A$, since in that case, $\hat{\delta}(w, v) \geq \hat{\delta}(A, v)$ for every $v \in V$. The subscript A is omitted when clear from the context.

Lemma 16. *Given a graph $\hat{G} = (V, \hat{E})$, a subset of vertices $U \subseteq V$, and an integer $s \in [1, n]$, Algorithm 2 computes a subset of vertices $A \subseteq V$, such that:*

- *the expected size of the set A is at most $2s \log n$; and*
- *for every vertex $w \in V$, we have $|\hat{C}_A(w) \cap U| \leq 4|U|/s$.*

Algorithm 2 Finding Centers for a Subset

```
1: function SUBSET-CENTERS(Ĝ, U, s)
2:     A ← ∅
3:     while there exists w ∈ V such that |Ĉ(w) ∩ U| > 4|U|/s do
4:         W ← {w ∈ V : |Ĉ(w) ∩ U| > 4|U|/s}
5:         Add each element of W to A with probability min(s/|W|, 1)
6:     return A
```

Algorithm 2 is a generalization of the algorithm CENTER in [17]; and Lemma 16 is a trivial extension of Theorem 3.1 in [17].[7]

[7] As noted in [17], it is possible to derandomize the center-selecting algorithm, and its running time is still polynomial.

Using a set of centers A, we define, for each $a \in A$, its *extended Voronoi cell* $\hat{D}_a \subseteq U$ as follows:

$$\hat{D}_a = \left(\bigcup \left\{\hat{C}(b) : b \in \hat{N}_2(a)\right\} \cup \hat{N}_2(a)\right) \cap U. \tag{2}$$

We define $C(w)$ and D_a similarly as $\hat{C}(w)$ and $\hat{D}_a$, but with respect to the graph G.

The following lemma is the base of the recursion. Its proof is similar to that of Lemma 3 in [11].

Lemma 17. $\bigcup_{a\in A} G[D_a]$ *covers every edge of* $G[U]$.

Proof. We prove that for every edge $\{u, v\}$ of $G[U]$, there is some $a \in A$, such that both u and v are in D_a. Let $\{u, v\}$ be any edge of $G[U]$. Without loss of generality, we assume $\delta(A, u) \leq \delta(A, v)$. We choose $a \in A$ such that $\delta(a, u) = \delta(A, u)$. If $\delta(a, u) \leq 1$, then both u and v are in $N_2(a) \cap U \subseteq D_a$. If $\delta(a, u) \geq 2$, let b be the vertex at distance 2 from a on a shortest a-to-u path in G. By the triangle inequality, we have $\delta(b, v) \leq \delta(b, u) + \delta(u, v) = \delta(b, u) + 1$. Since $\delta(b, u) = \delta(a, u) - 2$ and $\delta(a, u) = \delta(A, u) \leq \delta(A, v)$, we have $\delta(b, u) < \delta(A, u)$ and $\delta(b, v) < \delta(A, v)$. So both u and v are in $C(b) \cap U$, which is a subset of D_a since $b \in N_2(a)$. □

From Lemma 17, verifying that every edge of $G[U]$ is an edge of $\hat{G}[U]$ reduces to verifying that every edge of $G[D_a]$ is an edge of $\hat{G}[D_a]$ for every D_a. To see this, consider any edge $\{u, v\}$ of $G[U]$. There exists $a \in A$ such that $u, v \in D_a$. It is enough to verify that $\{u, v\}$ is an edge of $\hat{G}[D_a]$, hence an edge of $\hat{G}[U]$. This observation enables us to apply recursion on each D_a.

The main difficulty is: **How to obtain D_a efficiently?** If we compute D_a from its definition, we first need to compute $N_2(a)$, which takes too many queries since $N_2(a)$ may contain nodes outside U. Instead, a careful analysis shows that we can check whether $D_a = \hat{D}_a$ without even knowing $N_2(a)$, and $\hat{D}_a$ can be inferred from the graph $\hat{G}$ with no queries. This is shown in Lemma 18, which is the main novelty of the algorithmic design.

Lemma 18. *Assume that* $\hat{E} \subseteq E$. *If* $\delta(u, v) = \hat{\delta}(u, v)$ *for every pair* (u, v) *from* $\bigcup_{a\in A} \hat{N}_2(a) \times U$, *then* $D_a = \hat{D}_a$ *for all* $a \in A$.

Proof. The proof is delicate but elementary. For every $b \in \bigcup_{a\in A} \hat{N}_2(a)$, we have $\hat{C}(b) \cap U = C(b) \cap U$, because we have verified that $\hat{\delta}(b, u) = \delta(b, u)$ and $\hat{\delta}(A, u) = \delta(A, u)$ for every $u \in U$. Therefore, $\hat{D}_a$ can be rewritten as $\left(\bigcup \left\{C(b) : b \in \hat{N}_2(a)\right\} \cup \hat{N}_2(a)\right) \cap U$. Since $\hat{E} \subseteq E$, we have $\hat{N}_2(a) \subseteq N_2(a)$. Therefore $\hat{D}_a \subseteq D_a$.

On the other hand, we have $N_2(a) \cap U \subseteq \hat{N}_2(a) \cap U$, because we have verified that $\hat{\delta}(a, u) = \delta(a, u)$ for all u in $N_2(a) \cap U$. To prove $D_a \subseteq \hat{D}_a$, it only remains to show that, for any vertex $u \notin N_2(a)$ such that $u \in C(b) \cap U$ for some $b \in N_2(a)$, we have $u \in C(x) \cap U$ for some $x \in \hat{N}_2(a)$. We choose x to be the vertex

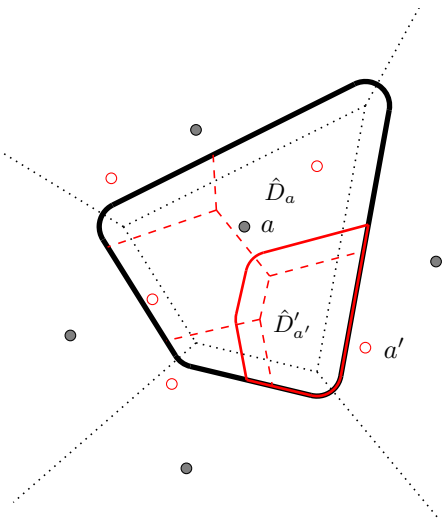

The solid points are top-level centers returned by SUBSET-CENTERS$(\hat{G}, V, s)$. The dotted lines indicate the partition of V into Voronoi cells by those centers. For a center a, expanding slightly its Voronoi cell results in $\hat{D}_a$ (the region inside the outer closed curve). On the second level of the recursive call for $\hat{D}_a$, the hollow points are the centers returned by SUBSET-CENTERS$(\hat{G}, \hat{D}_a, s)$. Observe that some of those centers lie outside $\hat{D}_a$. The dashed lines indicate the partition of $\hat{D}_a$ into Voronoi cells by those centers. Similarly, for a center a', expanding slightly its Voronoi cell results in $\hat{D}'_{a'}$ (the region inside the inner closed curve). Note that every $\hat{D}'_{a'}$ is inside $\hat{D}_a$.

Fig. 1: Two levels of recursive calls of VERIFY-SUBGRAPH$(\hat{G}, V)$

at distance 2 from a on a shortest a-to-u path in $\hat{G}$. By the assumption and the definition of x, we have: $\delta(x,u) = \hat{\delta}(x,u) = \hat{\delta}(a,u) - 2 = \delta(a,u) - 2$. By the triangle inequality, and using $b \in N_2(a)$ and $u \in C(b)$, we have: $\delta(a,u) \leq \delta(a,b) + \delta(b,u) \leq 2 + \delta(b,u) < 2 + \delta(A,u)$. Therefore $\delta(x,u) < \delta(A,u)$. Thus $u \in C(x) \cap U$. □

The recursive algorithm for non-edge verification is in Algorithm 3. It queries every $(u,v) \in \bigcup_{a \in A} \hat{N}_2(a) \times U$ and then recurses on each extended Voronoi cell $\hat{D}_a$. See Figure 1. It returns *yes* if and only if every query during the execution gives the right distance. The parameters n_0 and s are defined later. We assume that every edge of $\hat{G}$ has already been confirmed, i.e., $\hat{E} \subseteq E$. Correctness of the algorithm follows trivially from Lemmas 17 and 18.

Algorithm 3 Recursive Verification

```
1: procedure VERIFY-SUBGRAPH(Ĝ, U)
2:     if |U| > n_0 then
3:         A ← SUBSET-CENTERS(Ĝ, U, s)                ▷ Algorithm 2
4:         for a ∈ A do
5:             QUERY(N̂_2(a), U)
6:             VERIFY-SUBGRAPH(Ĝ, D̂_a)                ▷ D̂_a defined in Equation (2)
7:     else QUERY(U, U)
```

Next, we analysis the query complexity of VERIFY-SUBGRAPH($\hat{G}, V$). Define

$$k_0 = \left\lfloor \sqrt{\frac{\log n}{\log\left(\log n \cdot 32(\Delta^2+1)^2\right)}} \right\rfloor.$$

Let $s = n^{1/k_0}$ and $n_0 = \left(4(\Delta^2+1)\right)^{k_0}$ be the parameters in VERIFY-SUBGRAPH. Consider any recursive call when $|U| > n_0$. Let $A \subseteq V$ be the centers returned by SUBSET-CENTERS. By Lemma 16, $|A| \le 2s \log n$ and every $|\hat{C}(w) \cap U|$ is at most $4|U|/s$. Since the graph has maximum degree Δ, the size of every $\hat{D}_a$ is at most $(\Delta^2+1) \cdot \max(4|U|/s, 1)$. Therefore by induction, for any $1 \le k \le k_0+1$, any subset U on the k^{th} level of the recursion has size at most $t_k := n\left(4(\Delta^2+1)/s\right)^{k-1}$, where $t_{k_0+1} = n_0$. Hence the maximum level of the recursion is at most k_0+1.

First, consider the recursive calls with $|U| \le n_0$. There are at most $(2s \log n)^{k_0}$ such calls and each takes $|U|^2 \le \left(4(\Delta^2+1)\right)^{2k_0}$ queries. So their overall query complexity is at most $n \cdot \left(\log n \cdot 32(\Delta^2+1)^2\right)^{k_0} \le n^{1+1/k_0}$.

Next, consider the recursive calls with $|U| > n_0$ on the k^{th} level of the recursion for some fixed $k \in [1, k_0]$.[8] There are at most $(2s \log n)^{k-1}$ such calls and each takes at most $(\Delta^2+1)|A| \cdot |U|$ queries, where $|U| \le t_k$. So their overall query complexity is at most $n^{1+1/k_0}\left(\log n \cdot 8(\Delta^2+1)\right)^k$. Summing over k from 1 to k_0, the query complexity of all recursive calls with $|U| > n_0$ is at most $2 \cdot n^{1+1/k_0}\left(\log n \cdot 8(\Delta^2+1)\right)^{k_0} \le 2 \cdot n^{1+2/k_0}$.

Therefore, the overall query complexity is at most $3 \cdot n^{1+2/k_0}$, which is $n^{1+O\left(\sqrt{(\log\log n + \log\Delta)/\log n}\right)}$, as stated in Theorem 14.

Remark. *The recursive algorithm for non-edge verification in this section (as well as the one in Section 3.4) can be used for verification by itself. However, we only use its query complexity to provide guarantee for the greedy algorithm in Section 3.1, because the greedy algorithm is much simpler.*

3.4 Proof of Theorem 15

To show Theorem 15, we provide a recursive algorithm for non-edge verification of graphs of bounded treewidth with the query complexity in the theorem statement. The algorithm first computes (1/2)-balanced separator in $\hat{G}$ and use it to obtain a partition of V. Then it verifies the non-edges of G between different components in the partition. Finally, it recurses inside each component. But there is a catch because of the query oracle: by querying a pair (u, v), we would like to get back their distance in the recursive subgraph H, but instead the oracle returns their distance in the entire graph G. It could well be that a shortest u-to-v path in G goes through two nodes s_1 and s_2 in the separator where the segment between s_1 and s_2 is outside H.

[8] We note that there are no recursive calls on the $(k_0+1)^{\text{th}}$ level (i.e., last level) of the recursion with $|U| > n_0$.

As a warmup, we first provide an algorithm for the special case of chordal graphs, because the above issue does not arise when the graph is chordal.[9] We then extend the algorithm to graphs of bounded treewidth: To get around that issue, we formulate the recursive subproblem by augmenting H, adding virtual edges between vertices of the separator and giving them weight equal to their distance in G.

Verifying Chordal Graphs. We have a recursive algorithm to verify that every non-edge of $\hat{G}$ is a non-edge of G when G is a chordal graph (Algorithm 4). The algorithm returns *yes* if and only if every query during the execution gives the right distance.

Algorithm 4 Recursive Verification for Chordal Graphs

procedure VERIFY-CHORDAL($\hat{G}, U$)
  if $|U| > 4(\Delta + 1)$ **then**
    $S \leftarrow$ (1/2)-balanced clique separator of $\hat{G}[U]$ of size at most $\Delta + 1$
    QUERY(S, U) and obtain $N(S) \cap U$; QUERY($N(S) \cap U, U$)
    for every component C of $\hat{G}[U] \setminus S$ **do** VERIFY-CHORDAL($\hat{G}, C \cup S$)
  else QUERY(U, U)

Let $U \subseteq V$ represent the set of vertices for which we are currently verifying the induced subgraph. By Corollary 10, there is a (1/2)-balanced clique separator S of $\hat{G}[U]$.[10] We confirm the non-edges between different components of $\hat{G}[U] \setminus S$ by querying every pair $(u, v) \in (N(S) \cap U) \times U$. Then for each component C of $\hat{G}[U] \setminus S$, we recursively verify the non-edges inside $\hat{G}[C \cup S]$. The recursive call on the subset $C \cup S$ still use the global QUERY oracle. But because S is a clique in G, for any $u, v \in C \cup S$, any shortest u-to-v path in G stays inside $C \cup S$, so the value returned by QUERY(u, v) is the distance in $G[C \cup S]$. The following lemma shows correctness of Algorithm 4 and is a main idea of the algorithm.

Lemma 19. *Assume that $\hat{E} \subseteq E$. If $\delta(u, v) = \hat{\delta}(u, v)$ for every $(u, v) \in (N(S) \cap U) \times U$, then there is no edge in $G[U]$ between different components of $\hat{G}[U] \setminus S$.*

Proof. Let X and Y be any two different components in the partition of $\hat{G}[U] \setminus S$. Let x be any vertex in X and y be any vertex in Y. We show that $\{x, y\}$ is not an edge in $G[U]$. Let a (resp. b) be the vertex in $N(S)$ that is closest to x (resp. y) in $\hat{G}[U]$. Then $a \in X$ and $b \in Y$. Since $\hat{E} \subseteq E$, we have $\hat{N}(S) \subseteq N(S)$. It is then easy to see that $a, b \in (N(S) \cap U) \setminus S$. Without loss of generality, assume $\delta(a, x) \leq \delta(b, y)$.

[9] Since the separator is a clique, the shortest s_1-to-s_2 path is an edge, and thus belongs to H.

[10] We note that S can be computed in polynomial time and with no queries.

Since $(a, y) \in (N(S) \cap U) \times U$, we have $\delta(a, y) = \hat{\delta}(a, y)$. Any shortest path in $\hat{G}[U]$ from a to y goes through S, so

$$\hat{\delta}(a, y) \geq \hat{\delta}(a, S) + \hat{\delta}(S, y) = \hat{\delta}(a, S) + 1 + \hat{\delta}(b, y) = 2 + \hat{\delta}(b, y).$$

Since $(b, y) \in (N(S) \cap U) \times U$, we have $\hat{\delta}(b, y) = \delta(b, y)$. Therefore $\delta(a, y) \geq 2 + \delta(b, y) \geq 2 + \delta(a, x)$. By the triangle inequality, $\delta(x, y) \geq \delta(a, y) - \delta(a, x) \geq 2$. Thus $\{x, y\}$ is not an edge in $G[U]$. □

Since $\hat{G}[U]$ has maximum degree Δ and S has size at most $\Delta+1$, $\textsc{Query}(S, U)$ and $\textsc{Query}(N(S) \cap U, U)$ use $O(\Delta^2|U|)$ queries. Let $q(m)$ be the number of queries of $\textsc{Verify-Chordal}(\hat{G}, U)$ when $|U| = m$. We have

$$q(|U|) = O(\Delta^2|U|) + \sum_C q(|C| + |S|),$$

where $|U| = |S| + \sum_C |C|$ and S is a $(1/2)$-balanced separator. Hence $q(n) = O(\Delta^2 n \log n)$.

Remark. *We note that there are simpler algorithms for verifying chordal graphs, but the algorithm presented here conveys ideas that can be extended to verify graphs of bounded treewidth.*

Verifying Graphs of Bounded Treewidth. We extend Algorithm 4 to graphs of treewidth tw. The input specification is now the graph $\hat{G}$, a subset $U \subseteq V$, *plus* a set F of additional, new edges $\{u, v\}$ with weight $\delta(u, v)$. The set F is initially empty, and increases during the recursion. The algorithm verifies whether the metric of $(U, \hat{E}[U] \cup F[U])$ is identical to that of $(U, E[U] \cup F[U])$. Instead of S being a clique, now S is an existing bag of some tree decomposition of width tw (see Lemma 8). Lemma 19 still holds. We create new edges $\{u, v\}$ with weight $w(u, v) := \delta(u, v)$ for all pairs $\{u, v\} \subseteq S$, and we add them to the set F. For each connected component C of $\hat{G}[U] \setminus S$, we make a recursive call for the vertex set $C \cup S$ and the updated set F of weighted edge. Every subgraph in the recursive call has treewidth at most tw, since the new edges are added inside S. This concludes the description and correctness of the algorithm.

For the query complexity, we need to bound the size of the neighborhood $N(S)$ of S: it is with respect to the subgraph $E[U] \cup F[U]$, so the vertex degree is no longer bounded by Δ. However, for any vertex v, the number of weighted edges adjacent to v is bounded by the maximum bag size times the number of bags S containing v that have been used as separators in the recursive calls. Since the graph has treewidth tw, every bag has size at most $\text{tw} + 1$. Since all separators are $(1/2)$-balanced, the recursion has depth $O(\log n)$, so v belongs to $O(\log n)$ such bags. Therefore, the degree of v is $O(\Delta + \text{tw} \log n)$. The overall query complexity is $O(\Delta(\Delta + \text{tw} \log n) n \log n)$.

Thus we proved Theorem 15.

Algorithm 5 Greedy Reconstruction

procedure RECONSTRUCT(V)
  $u_0 \leftarrow$ an arbitrary vertex
  for $u \in V \setminus \{u_0\}$ **do** QUERY(u, u_0) to get a shortest u-to-u_0 path
  $X \leftarrow$ the union of the above paths
  $Y \leftarrow \emptyset$
  while $X \cup Y$ does not cover all vertex pairs **do**
    choose (u, v) that maximizes $|S^X_{u,v} \setminus Y|$ ▷ $S^X_{u,v}$ defined in Equation (3)
    Query(u, v) to get a shortest u-to-v path
    if $\delta_G(u, v) = \delta_X(u, v)$ **then**
      $Y \leftarrow Y \cup S^X_{u,v}$
    else
      let e be some edge of the above u-to-v path that does not belong to X
      $X \leftarrow X \cup \{e\}$
  return X

4 Proof of Theorems 3

The algorithm (Algorithm 5) constructs an increasing set X of edges so that in the end $X = E$. At any time, the candidate graph is X.[11] Initially, X is the union of the shortest paths given as answers by $n - 1$ queries, so that X is a connected subgraph spanning V. At each subsequent step, the algorithm makes a query that leads either to the confirmation of many non-edges of G, or to the discovery of an edge of G.

Formally, we define, for every pair $(u, v) \in V^2$,

$$S^X_{u,v} = \big\{\{a, b\} \text{ is an non-edge of } X : \delta_X(u, a) + \delta_X(b, v) + 1 < \delta_X(u, v)\big\}. \quad (3)$$

This is similar to $S_{u,v}$ defined in Equation (1). From Lemma 12, $S^X_{u,v}$ contains the pairs that can be confirmed as non-edges of G if $\delta_G(u, v) = \delta_X(u, v)$. At each step, the algorithm queries a pair (u, v) that maximizes the size of the set $S^X_{u,v} \setminus Y$. As a consequence, either all pairs in $S^X_{u,v} \setminus Y$ are confirmed as non-edges of G, or $\delta_G(u, v) \neq \delta_X(u, v)$, and in that case, the query reveals an edge along a shortest u-to-v path in G that is not in X; we then add this edge to X.

To see the correctness, we note that the algorithm maintains the invariant that all pairs in X are confirmed edges of G, and that all pairs in Y are confirmed non-edges of G. Thus when $X \cup Y$ covers all vertex pairs, we have $X = E$.

For the query complexity, first, consider the queries that lead to $\delta_G(u, v) \neq \delta_X(u, v)$. For each such query, an edge is added to X. This can happen at most $|E| \leq \Delta n$ times, because the graph has maximum degree Δ.

Second, consider the queries that lead to $\delta_G(u, v) = \delta_X(u, v)$. Define R to be the set of vertex pairs that are not in $X \cup Y$. We analyze the size of R. For each such query, the size of R decreases by $|S^X_{u,v} \setminus Y|$. To lower bound $|S^X_{u,v} \setminus Y|$, we consider the problem of non-edge verification using a distance oracle on the input

[11] We identify X with the subgraph induced by the edges of X.

graph X, and let T be an (unknown) optimal set of queries. By Theorem 14, $|T|$ is at most $f(n,\Delta) = n^{1+O\left(\sqrt{(\log\log n+\log\Delta)/\log n}\right)}$. By Lemma 13, the sets $S^X_{u,v}$ for all pairs $(u,v) \in T$ together cover $R \cup Y$, hence R. Therefore, at least one of these pairs satisfies $|S^X_{u,v} \setminus Y| \geq |R|/|T|$. Initially, $|R| \leq n(n-1)/2$, and right before the last query, $|R| \geq 1$, thus the number of queries with $\delta_G(u,v) = \delta_X(u,v)$ is $O(\log n) \cdot f(n,\Delta)$.

Therefore, the overall query complexity is at most $(n-1) + \Delta n + O(\log n) \cdot f(n,\Delta)$. Thus we obtained the same query bound as in the first statement of Theorem 1. To prove the query bound for graphs of treewidth tw as in the second statement, the analysis is identical as above, except that we use Theorem 15 instead of Theorem 14 to obtain $f(n,\Delta)$.

Remark. *Note that the above proof depends crucially on the fact that $f(n,\Delta)$ is a uniform bound on the number of distance queries for non-edge verification of* any *n-vertex graph of maximum degree Δ. Thus, even though the graph X changes during the course of the algorithm because of queries (u,v) such that $\delta_G(u,v) \neq \delta_X(u,v)$, each query for which the distance in G and the current X are equal confirms $1/f(n,\Delta)$ fraction of non-edges.*

5 Proof of Theorem 4

The algorithm for Theorem 4 uses a clique separator to partition the graph into balanced subgraphs, and then recursively reconstructs each subgraph. The main difficulty is to compute the partition. The partition algorithm and its analysis are the main novelty in this section, see Section 5.1. In what follows, the set U represents the set of vertices for which we are currently reconstructing the induced subgraph during the recursion.

Definition 20. *A subset of vertices $U \subseteq V$ is said to be* self-contained *if, for every pair of vertices $(x,y) \in U^2$, any shortest path in G between x and y goes through nodes only in U.*

The set U during the recursion is always *self-contained*, because every separator is a clique.

5.1 Subroutine: Computing the Partition

Let U be a self-contained subset of V. Let S be a subset of U. We want to compute the partition of $G[U] \setminus S$ into connected components. Let $W = (N(S) \cap U) \setminus S$. For every $a \in W$, define $B(a)$ as the *cluster* at a:

$$B(a) = \{x \in U \setminus S \mid \delta(a,x) \leq \delta(S,x)\}. \tag{4}$$

Since U is self-contained, every $x \in U \setminus S$ belongs to some cluster $B(a)$. However, the clusters may have overlaps. The algorithm (see Algorithm 6) successively merges two clusters with overlaps. See Figure 2.

Algorithm 6 Computing the Partition

1: **function** PARTITION(U, S)
2: QUERY(S, U) and obtain $N(S) \cap U$; QUERY$(N(S) \cap U, U)$
3: $W \leftarrow (N(S) \cap U) \setminus S$
4: $\mathcal{B} \leftarrow \{B(a) \mid a \in W\}$ ▷ $B(a)$ defined in Equation (4)
5: **while** $\exists\ B_1, B_2 \in \mathcal{B}$ s.t. $B_1 \cap B_2 \neq \emptyset$ **do** merge B_1 and B_2 in $\mathcal{B}$
6: **return** $\mathcal{B}$

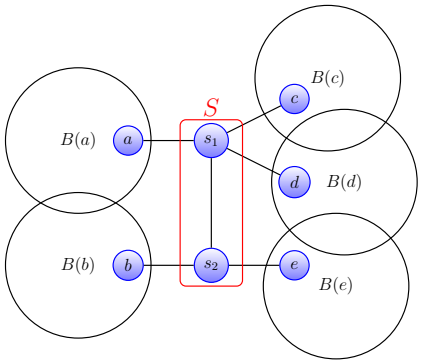

In the example, $S = \{s_1, s_2\}$ and $W = \{a, b, c, d, e\}$. The clusters $B(a)$, $B(b)$, $B(c)$, $B(d)$, $B(e)$ are indicated by the balls. Using their overlaps, the algorithm produces the partition $\mathcal{B} = \{B(a) \cup B(b)\,,\, B(c) \cup B(d) \cup B(e)\}$.

Fig. 2: Example of the Partition

Lemma 21. *Algorithm* PARTITION *uses* $O(\Delta|S| \cdot |U|)$ *queries and outputs the partition of* $G[U] \setminus S$ *into connected components.*

The query complexity of the algorithm is $O(|N(S)| \cdot |U|) = O(\Delta|S| \cdot |U|)$. Lemma 21 then follows directly from Lemmas 22 and 23.

Lemma 22. *Let* C *be a connected component in* $G[U] \setminus S$*. Then* $C \subseteq B$ *for some set* B *in the output of the algorithm.*

Proof. Let A be the set of vertices in $C \cap W$. Since U is self-contained, for every vertex $x \in C$, there exists some $a \in A$ such that $x \in B(a)$. Thus we only need to prove that all sets $\{B(a) : a \in A\}$ are eventually merged in our algorithm.

Define a weighed graph H whose vertex set is A, and such that for every $(a, b) \in A^2$, there is an edge (a, b) in H with weight $w(a, b)$, which is defined as the distance between a and b in $G[C]$[12]. To show that all sets $\{B(a) : a \in A\}$ are eventually merged, we use an inductive proof that is in the same order that Prim's algorithm would construct a minimum spanning tree on H. Recall that Prim's algorithm initializes a tree $\mathcal{T}$ with a single vertex, chosen arbitrarily from A. Then it repeatedly chooses an edge $(a, b) \in \mathcal{T} \times (A \setminus \mathcal{T})$ with minimum weight and add this edge to $\mathcal{T}$. We will show that if an edge (a, b) is added to $\mathcal{T}$, then $B(a)$ and $B(b)$ are merged in our algorithm. Since Prim's algorithm finishes by

[12] This distance may be larger than $\delta(a, b)$, the distance between a and b in G.

providing a spanning tree including every $a \in A$, we thus proved that all sets $B(a)$ for $a \in A$ are merged in our algorithm.

Suppose that the i unions corresponding to the first i edges chosen by Prim's algorithm have been performed already, for $i \geq 0$. Let $\mathcal{T}$ be the tree in H after adding the first i edges.[13] Let (a, b) be the $(i+1)^{\text{th}}$ edge chosen by Prim's algorithm. Thus $a \in \mathcal{T}$, $b \in A \setminus \mathcal{T}$, and $w(a, b)$ is minimized. Consider a shortest path $p_1, \ldots, p_k$ in $G[C]$ between a and b. Let $z = p_{\lceil k/2 \rceil}$ be the mid-point vertex of the path. We show that both $B(a)$ and $B(b)$ contain z, thus $B(a)$ and $B(b)$ are merged in our algorithm. It is easy to see that $p_1, \ldots, p_{\lceil k/2 \rceil}$ and $p_{\lceil k/2 \rceil}, \ldots, p_k$ are shortest paths in G. Thus $\delta(a, z) = \lceil k/2 \rceil - 1$ and $\delta(b, z) = \lfloor k/2 \rfloor$. So we have $\delta(a, z) \leq \delta(b, z) \leq \delta(a, z) + 1$. To show $z \in B(a)$ and $z \in B(b)$, we only need to show that $\delta(b, z) \leq \delta(S, z)$. Choose the vertex $s \in S$ that minimizes $\delta(s, z)$ and consider a shortest z-to-s path P. Let c be the neighbor of s on P, and let P' be the shortest z-to-c path. We note that $c \in A$ and P' is in $G[C]$. Since $\delta(S, z) = \delta(s, z) = \delta(c, z) + 1$, we only need to show that $\delta(b, z) \leq \delta(c, z) + 1$. There are 2 cases:

Case 1: $c \in A \setminus \mathcal{T}$. Then the concatenation of $p_1, \ldots, p_{\lceil k/2 \rceil}$ and P' gives a path in $G[C]$ between a and c of length $\delta(a, z) + \delta(c, z)$, which is at least $w(a, c)$ by the definition of the weight. From the choice of (a, b), $w(a, c) \geq w(a, b) = \delta(a, z) + \delta(b, z)$. So we have $\delta(b, z) \leq \delta(c, z)$.

Case 2: $c \in \mathcal{T}$. Similarly, the concatenation of $p_k, p_{k-1}, \ldots, p_{\lceil k/2 \rceil}$ and P' gives a path in $G[C]$ between b and c of length $\delta(b, z) + \delta(c, z)$, which is at least $w(b, c)$ by the definition of the weight. From the choice of (a, b), $w(b, c) \geq w(a, b) = \delta(a, z) + \delta(b, z)$. So we have $\delta(a, z) \leq \delta(c, z)$. Thus $\delta(b, z) \leq \delta(a, z) + 1 \leq \delta(c, z) + 1$. □

Lemma 23. *Let B be a set in the output of the algorithm. Then $B \subseteq C$ for some connected component C in $G[U] \setminus S$.*

Proof. First we show that for every $a \in W$ and every $x \in B(a)$, a and x belong to the same component in $G[U] \setminus S$. Suppose there exists some $x \in B(a)$, such that x and a belong to different components in $G[U] \setminus S$. Any shortest path from a to x must pass through the separator S, so we have $\delta(a, x) \geq \delta(a, S) + \delta(S, x) = 1 + \delta(S, x)$. Contradiction with $x \in B(a)$.

Next we prove an invariant on $\mathcal{B}$ during the **while** loop (Line 5): *Every set $B \in \mathcal{B}$ is a subset of some component of $G[U] \setminus S$.* This invariant holds before the **while** loop starts. Suppose the invariant holds before the i^{th} iteration of the **while** loop, and in this iteration $B_1, B_2 \in \mathcal{B}$ get merged. Since $B_1 \cap B_2 \neq \emptyset$, there exists $z \in B_1 \cap B_2$. All nodes in B_1 (resp. in B_2) are in the same component as z. Thus all nodes in $B_1 \cup B_2$ are in the same component as z. By induction, the invariant holds when the **while** loop terminates.

Thus we complete the proof. □

[13] For the base case ($i = 0$), $\mathcal{T}$ contains a single vertex and no union operation is performed.

Algorithm 7 Finding a Shortest Path

```
1: function SHORTEST-PATH(U, a, b)
2:     if δ(a, b) > 1 then
3:         QUERY(a, U); QUERY(b, U)
4:         T ← {v ∈ U | δ(v, a) + δ(v, b) = δ(a, b)}
5:         l ← ⌊δ(a, b)/2⌋
6:         c ← an arbitrary node in T such that δ(c, a) = ℓ
7:         U_1 ← {v ∈ T | δ(v, a) < ℓ}
8:         U_2 ← {v ∈ T | δ(v, a) > ℓ}
9:         P_1 ← SHORTEST-PATH(U_1, a, c)
10:        P_2 ← SHORTEST-PATH(U_2, c, b)
11:        return the concatenation of P_1 and P_2
12:    else
13:        return the path of a single edge (a, b)
```

5.2 Subroutine: Computing a Shortest Path

Given a self-contained subset of vertices $U \subseteq V$ and two vertices $a, b \in U$, Algorithm 7 computes a shortest path between a and b by divide-and-conquer. The query complexity is $O(|U| \log |U|)$. See Appendix A.1 of [11] for the analysis of the algorithm.

5.3 Algorithm and Analysis

The reconstruction algorithm is in Algorithm 8. To find a balanced separator, we use ideas from [11]: the algorithm computes a vertex that is on many shortest paths in the sampling, and grows a clique including this vertex. The constants n_0, C_1, and $0 < \beta < 1$ are defined later.

Lemma 24. RECONSTRUCT-CHORDAL(U) *indeed returns the edge set of* $G[U]$.

Proof. By Lemma 21, $U_1, \ldots, U_\ell$ are the connected components in $G[U] \setminus K$. There cannot be edges between different U_i and U_j. Thus every edge of $G[U]$ belongs to some $G[U_i \cup K]$. So the edge set of $G[U]$ is the union of the edge sets of $G[U_i \cup K]$ over i. Hence correctness follows by induction. □

The rest of this section is to analyze the query complexity. We set the constants $n_0 = 2^{\Delta+2}(\Delta+1)^2$; $\beta = \max\left(1 - 1/(\Delta \cdot 2^{\Delta+1}), \sqrt{1 - 1/(4(\Delta+1))}\right)$; and $C_1 = 256(\Delta+1)^2$. The key is the following lemma.

Lemma 25. *In every* ***repeat*** *loop of* BALANCED-SEPARATOR*, a* β*-balanced separator is found with probability at least* $2/3$.

We defer the proof of Lemma 25 to Section 5.4 and show in the rest of this section how Lemma 25 implies the query complexity stated in Theorem 4.

First we analyze the query complexity of BALANCED-SEPARATOR. Computing $C_1 \log |U|$ shortest paths takes $O(\Delta^2 |U| \log^2 |U|)$ queries, since a shortest

Algorithm 8 Reconstruction of Chordal Graphs

procedure RECONSTRUCT-CHORDAL(U)
  if $|U| > n_0$ **then**
    $K \leftarrow$BALANCED-SEPARATOR(U)
    $(U_1, \ldots, U_\ell) \leftarrow$PARTITION($U, K$) ▷ See Algorithm 6
    return $\bigcup_i$RECONSTRUCT-CHORDAL($U_i \cup K$)
  else
    reconstruct $G[U]$ by QUERY(U, U)
function BALANCED-SEPARATOR(U) ▷ finds a β-balanced separator of $G[U]$
  repeat
    for $i \leftarrow 1$ **to** $C_1 \log |U|$ **do**
      $(a_i, b_i) \leftarrow$ a pair of uniformly random nodes from U
      $P_i \leftarrow$ SHORTEST-PATH(a_i, b_i, U) ▷ see Section 5.2
    $x \leftarrow$ the node in U with the most occurrences among all P_i's
    QUERY(x, U) and obtain $N(x)$
    QUERY($N(x), N(x)$) and obtain all cliques containing x
    for every clique K containing x **do**
      $(U_1, \ldots, U_\ell) \leftarrow$ PARTITION(U, K) ▷ See Algorithm 6
      if $\max_i |U_i| < \beta|U|$ **then return** K
  until a balanced separator is found

path between two given nodes can be computed using $O(|U| \log |U|)$ queries (see Section 5.2). We note that the neighborhood $N(x)$ of x has size at most $\Delta + 1$, and there are at most 2^Δ cliques containing x. By Lemma 21, PARTITION(U, K) takes $O(\Delta|K| \cdot |U|)$ queries, where $|K| \leq \Delta + 1$. Therefore every **repeat** loop in BALANCED-SEPARATOR takes $O\left(\Delta^2|U|(2^\Delta + \log^2 |U|)\right)$ queries. By Lemma 25, the expected number of **repeat** loops is constant. So the query complexity of BALANCED-SEPARATOR is $O\left(\Delta^2|U|(2^\Delta + \log^2 |U|)\right)$.

Next, we analyze the query complexity of RECONSTRUCT-CHORDAL(U). Let $q(m)$ be the number of queries when $|U| = m$. We have

$$q(|U|) = O\left(\Delta^2|U|(2^\Delta + \log^2 |U|)\right) + \sum_i q(|U_i| + |K|),$$

where $|U| = |K| + \sum_i |U_i|$ and K is a β-balanced separator of size at most $\Delta + 1$. Hence $q(n) = O\left(\Delta^2 n(2^\Delta + \log^2 n) \log_{\frac{1}{\beta}} n\right) = O\left(\Delta^3 2^\Delta \cdot n(2^\Delta + \log^2 n) \log n\right)$.

5.4 Proof of Lemma 25

First, we need Lemmas 26 and 27.

Lemma 26. *For $v \in U$, let p_v denote the fraction of pairs $(a, b) \in U^2$ such that v is on some shortest path between a and b. Then $\max_v p_v \geq 1/(2(\Delta + 1))$.*

Proof. By Corollary 10, there is some clique separator S of size at most $\Delta + 1$ such that every connected component in $G[U] \setminus S$ has size at most $|U|/2$. Notice

that for any pair of vertices a, b from different components, any shortest a-to-b path must go by some node in S. The number of such pairs is at least $|U|^2/2$. By Pigeonhole Principle, there exists some $z \in S$, such that for at least $1/|S| \geq 1/(\Delta+1)$ fraction of these pairs, their shortest paths go by z. Thus $p_z \geq 1/(2(\Delta+1))$. □

Lemma 27 (slightly adapted from [11]). *For ever vertex $v \in U$, let $\hat{p}_v$ denote the fraction of pairs (a_i, b_i) among $C_1 \log |U|$ uniformly and independently random pairs of U^2 such that v is on some shortest path between a_i and b_i. Let $x = \arg\max \hat{p}_x$. Then with probability at least $2/3$, we have $p_x > (\max_v p_v)/2$.*

Now we prove Lemma 25. By Lemma 9, there is a tree decomposition T of $G[U]$ such that every bag of T is a unique maximal clique of $G[U]$. Let x be the node computed on Line 13 of Algorithm 8. Let T_x be the subtree of T induced by the bags containing x. Define F to be the forest after removing T_x from T. For any subgraph H of T, define $V(H) \subseteq U$ to be the set of vertices that appear in at least one bag of H.

Case 1: There exists some connected component T' in F with $(1-\beta)|U| \leq |V(T')| \leq \beta|U|$. Consider the edge (K_1, K_2) in T such that $K_1 \in T_x$ and $K_2 \in T'$. $K_1 \cap K_2$ is a β-balanced separator, since $V(T')$ is a component in $G[U] \setminus (K_1 \cap K_2)$. Thus $K_1 \supseteq K_1 \cap K_2$ is also a β-balanced separator. Observe that $x \in K_1$, so K_1 is one of the cliques checked on Line 16. The algorithm succeeds by finding a β-balanced separator.

Case 2: There exists some connected component T' in F with $|V(T')| > \beta|U|$. The algorithm then fails to find a β-balanced separator. We bound the probability of this case by at most $1/3$. Again let (K_1, K_2) be the edge in T such that $K_1 \in T_x$ and $K_2 \in T'$. For any vertices $u, v \in V(T')$, any shortest u-to-v path cannot go by x. Since there are at least β^2 fraction of such pairs in U^2, we have $p_x \leq 1 - \beta^2$, which is at most $1/(4(\Delta+1))$ by the definition of β. This happens with probability at most $1/3$ by Lemmas 26 and 27.

We argue that the two cases above are exhaustive. Suppose, for the sake of contradiction, that every component T' in F is such that $|V(T')| < (1-\beta)|U|$. The number of components in F is at most $\Delta \cdot 2^{\Delta}$, because every component has a bag that contains a neighbor of x, and all bags are unique. So $|V(F)| < \Delta \cdot 2^{\Delta} \cdot (1-\beta)|U|$, which is at most $|U|/2$ by the definition of β. On the other hand, every node $v \in U \setminus N(x)$ is covered by some clique in F, so $|V(F)| \geq |U| - (\Delta+1)$, which is greater than $|U|/2$ since $|U| > n_0$. Contradiction.

Thus we complete the proof of Lemma 25.

6 Lower Bounds

6.1 Lower Bound for Graphs of Unbounded Degree

Reconstruction of graphs of unbounded degree using a distance oracle requires $\Omega(n^2)$ queries [14]. This lower bound can be easily extended to verification

or/and to the shortest path oracle model as follows. Consider the graph G of vertices $v_1, \ldots, v_n$, which contains a star: it has an edge $\{v_1, v_i\}$ for every $2 \le i \le n$. G may or may not contain one additional edge $\{v_i, v_j\}$ for $2 \le i, j \le n$. (In the verification version of the problem, the star graph is given as $\hat{G}$.) To detect if G contains such an edge $\{v_i, v_j\}$ for $2 \le i, j \le n$, we need to perform $\Omega(n^2)$ distance or shortest path queries.

6.2 Lower Bound for Reconstruction of Bounded Degree Graphs

We assume that $n = 3t - 1$ where $t = 2^k$ for some integer k. (The general case is similar.) Consider a family $\mathcal{G}$ of graphs G as follows: the vertex set is $\{v_1, \ldots, v_n\}$; the first $2t - 1$ vertices form a complete binary tree of height k (with leaves $v_t, \ldots, v_{2t-1}$); the other vertices $v_{2t}, \ldots, v_{3t-1}$ induce an arbitrary subgraph of maximum degree $\Delta - 1$; there is an edge between v_i and v_{i+t} for every $i \in [t, 2t - 1]$ and there are no other edges. Then every vertex in G has degree at most Δ, and the diameter of the graph is at most $2k+2$. Every distance query returns a number between 1 and $2k+2 = O(\log n)$, so it gives $O(\log \log n)$ bits of information. From information theory, the number of queries is at least the logarithm of the number of graphs in $\mathcal{G}$ divided by the maximum number of bits of information per query. The number of graphs in $\mathcal{G}$ is the number of graphs of size t and of maximum degree $\Delta - 1$, which is $\Omega\left(n^{\Omega(\Delta n)}\right)$ when $\Delta = o(\sqrt{n})$ (see [12]). Therefore, we have a query lower bound of

$$\frac{\log\left(\Omega\left(n^{\Omega(\Delta n)}\right)\right)}{O(\log \log n)} = \Omega\left(\frac{\Delta n \log n}{\log \log n}\right).$$

Acknowledgments. We thank Uri Zwick for Theorem 5. We thank Fabrice Benhamouda, Mathias Bæk Tejs Knudsen, and Mikkel Thorup for discussions.

References

1. Achlioptas, D., Clauset, A., Kempe, D., Moore, C.: On the bias of traceroute sampling: or, power-law degree distributions in regular graphs. Journal of the ACM (JACM) 56(4), 21 (2009)
2. Beerliova, Z., Eberhard, F., Erlebach, T., Hall, E., Hoffmann, M., Ram, L.S.: Network discovery and verification. In: WG. pp. 127–138. Springer (2005)
3. Blair, J.R.S., Peyton, B.: An introduction to chordal graphs and clique trees. In: Graph theory and sparse matrix computation, pp. 1–29. Springer (1993)
4. Castro, R., Coates, M., Liang, G., Nowak, R., Yu, B.: Network tomography: recent developments. Statistical Science 19, 499–517 (2004)
5. Chung, F., Garrett, M., Graham, R., Shallcross, D.: Distance realization problems with applications to internet tomography. Journal of Computer and System Sciences 63, 432–448 (2001)
6. Dall'Asta, L., Alvarez-Hamelin, I., Barrat, A., Vázquez, A., Vespignani, A.: Exploring networks with traceroute-like probes: Theory and simulations. Theoretical Computer Science 355(1), 6–24 (2006)

7. Erlebach, T., Hall, A., Hoffmann, M., Mihal'ák, M.: Network discovery and verification with distance queries. Algorithms and Complexity pp. 69–80 (2006)
8. Hein, J.J.: An optimal algorithm to reconstruct trees from additive distance data. Bulletin of Mathematical Biology 51(5), 597–603 (1989)
9. Johnson, D.S.: Approximation algorithms for combinatorial problems. Journal of computer and system sciences 9(3), 256–278 (1974)
10. King, V., Zhang, L., Zhou, Y.: On the complexity of distance-based evolutionary tree reconstruction. In: SODA. pp. 444–453. SIAM (2003)
11. Mathieu, C., Zhou, H.: Graph reconstruction via distance oracles. In: ICALP (1). pp. 733–744. Springer (2013)
12. McKay, B.D., Wormald, N.C.: Asymptotic enumeration by degree sequence of graphs with degrees $o(n^{1/2})$. Combinatorica 11(4), 369–382 (1991)
13. Reed, B.A.: Algorithmic aspects of tree width. In: Recent advances in algorithms and combinatorics, pp. 85–107. Springer (2003)
14. Reyzin, L., Srivastava, N.: Learning and verifying graphs using queries with a focus on edge counting. In: Algorithmic Learning Theory. pp. 285–297. Springer (2007)
15. Reyzin, L., Srivastava, N.: On the longest path algorithm for reconstructing trees from distance matrices. Information processing letters 101(3), 98–100 (2007)
16. Tarissan, F., Latapy, M., Prieur, C.: Efficient measurement of complex networks using link queries. In: INFOCOM Workshops. pp. 254–259. IEEE (2009)
17. Thorup, M., Zwick, U.: Compact routing schemes. In: Symposium on Parallel Algorithms and Architectures. pp. 1–10. ACM (2001)